REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
06-04-2011	Technical Paper	MAR 2011 - APR 2011

4. TITLE AND SUBTITLE

Performance of BGP Among Mobile Military Networks

5a. CONTRACT NUMBER

FA8720-05-C-0002

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Glenn Carl and Scott Arbiv

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

1777 N Kent St, Suite 9030
Arlington, VA 22209

10. SPONSOR/MONITOR'S ACRONYM(S)

OSD-ATL, DDR&E

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

As tactical military networks deploy new IP-capable radios (e.g., JTRS), it is expected that network connectivity wiII increase both within and among Service and Coalition networking domains (e.g., Army, Navy, NATO). However, unmanaged use of this Increased connectivity may result in violation of some mission-critical operational limits. One mechanism known to (coarsely) manage large-scale IP networks is BGP routing policy. As such, this paper begins the study of BGP's applicability to manage an evolution of the GIG In which IP-based tactical radios proliferate (i.e., the Future GIG). To this end, this paper first presents a modification to BGP that allows for dynamic management of its peering sessions to accommodate network node mobility. Next, since it is known that BGP can have performance issues (e.g., slow convergence), this paper uses network emulation^1 to perform a performance assessment of our modified BGP protocol. Specifically, for ten independent realizations of a mobile wireless networking model, our modified BGP protocol is evaluated with respect to its generated protocol overhead, its ability to develop valid routes to destinations (e.g., reachability), and its influence on network's outage events. Furthermore, our modified BGP protocol is compared to the OSPF and OSPF-MDR routing protocols for mobile networks with increasing number of nodes. Our results show that modified BGP's overhead growth is significantly higher than both OSPF variations, but the network's average reachability and median outage times are similar. To decrease the significant overhead seen for BGP, this paper also presents a second modification to the BGP protocol based on the use of a connected dominating set (CDS) backbone. Here, the emulation results show that the use of a CDS backbone can significantly decrease BGP's overhead with little impact on the network's average reachability and median duration of its outage events. Lastly, it was found the maximum network outage times for both modified BGP protocols were tens of seconds greater than that experienced by either OSPF-based routing protocols. However, the occurrence rate of long network outage events (e.g., greater than 60 secs) was also seen to be infrequent for networks whose average shortest paths were less than 3 hops.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
U					Zach Sweet
a. REPORT	b. ABSTRACT	c. THIS PAGE	SAR	9	19b. TELEPHONE NUMBER *(include area code)*
U	U	U			781-981-5997

MS-56761

Performance of BGP Among Mobile Military Networks

Glenn Carl
Wideband Tactical Networking
MIT Lincoln Laboratory
Lexington, MA

Scott Arbiv
Wideband Tactical Networking
MIT Lincoln Laboratory
Lexington, MA

Abstract*—As tactical military networks deploy new IP-capable radios (e.g., JTRS), it is expected that network connectivity will increase both within and among Service and Coalition networking domains (e.g., Army, Navy, NATO). However, unmanaged use of this increased connectivity may result in violation of some mission-critical operational limits. One mechanism known to (coarsely) manage large-scale IP networks is BGP routing policy. As such, this paper begins the study of BGP's applicability to manage an evolution of the GIG in which IP-based tactical radios proliferate (i.e., the Future GIG). To this end, this paper first presents a modification to BGP that allows for dynamic management of its peering sessions to accommodate network node mobility. Next, since it is known that BGP can have performance issues (e.g., slow convergence), this paper uses network emulation[1] to perform a performance assessment of our modified BGP protocol. Specifically, for ten independent realizations of a mobile wireless networking model, our modified BGP protocol is evaluated with respect to its generated protocol overhead, its ability to develop valid routes to destinations (e.g., reachability), and its influence on network's outage events. Furthermore, our modified BGP protocol is compared to the OSPF and OSPF-MDR routing protocols for mobile networks with increasing number of nodes. Our results show that our modified BGP's overhead growth is significantly higher than both OSPF variations, but the network's average reachability and median outage times are similar. To decrease the significant overhead seen for BGP, this paper also presents a second modification to the BGP protocol based on the use of a connected dominating set (CDS) backbone. Here, the emulation results show that the use of a CDS backbone can significantly decrease BGP's overhead with little impact on the network's average reachability and median duration of its outage events. Lastly, it was found the maximum network outage times for both modified BGP protocols were tens of seconds greater than that experienced by either OSPF-based routing protocols. However, the occurrence rate of long network outage events (e.g., greater than 60 secs) was also seen to be infrequent for networks whose average shortest paths were less than 3 hops.

I. Introduction

As mobile tactical military networks deploy new interoperable communication waveforms (e.g., JTRS IP-capable radios), opportunities will be created to increase local cooperation between Service and Coalition networking domains (e.g., Army, Navy, and NATO). Specifically, IP-capable software-defined-radios deployed to tactical military networks are expected to increase local interconnectivity between many heterogenous units (e.g., naval platforms, strike aircraft, ISR assets, and ground forces) from different Service and Coalition organizations (e.g., Navy, Army, NATO) [1]. This cooperative networking arrangement is often claimed to provide several network benefits (e.g., increased any-to-any connectivity, better load balancing across bandwidth-limited wireless links) as well as providing the opportunity to distribute many networking services and applications across a shared IP networking cloud. However, unmanaged cooperation among these new net-centric platforms may result in violation of security and/or operational limits for some mission-critical network nodes. To manage this cooperation, IP routing policy can be implemented at the Services' and Coalitions' network boundaries to vary the amounts of inter-networking collaboration (via enforcement of performance and/or resource-utilization constraints). For example, a Navy networking domain may advertise IP routes to an Army domain to implicitly convey the Navy's offer to use its network resources to transit Army traffic (e.g., GIG reachback via ship-based SATCOM terminals).

To bring flexible routing policy capabilities into this Future GIG architecture, the Internet's de-facto policy-based routing protocol BGP [2] is considered. BGP is the most widely used policy-based routing protocol to date. Explicit MANET routing protocols (e.g., AODV, DSR, LAR [3]) are not considered policy-based since they appear to have the implicit constraint that all nodes within the routing domain use a common routing policy (e.g., shortest-path routing). It is conjectured that routing protocols that can only support a single routing policy will not have enough capability or flexibility to cooperatively manage multiple, but separately administered networking domains tasked with supporting multiple military missions. Furthermore, most emerging tactical networks (e.g., JTRS, WIN-T) have already specified the use of BGP to internetwork themselves to the GIG.

The ability of BGP to operate effectively among tactical networks continues to be questioned by the DoD networking community. The BGP routing protocol has been observed to exhibit high communication overhead [4] and slow routing

This work is sponsored by the Department of Defense under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by the United States Government.

[1]Network emulation is considered a testing practice which uses real hardware, protocols, and control software to model in real-time the behavior of an experimental network under test.

convergence [5] while operating in the Internet. Although the Internet continues to grow and operate in spite of such BGP limitations, this may not be the case for the Future GIG. The Future GIG is envisioned to be dominated by mobile nodes having lower bandwidth and more volatile RF links (whereas the Internet's generally contains fixed, robust, higher-rate optical links). A performance assessment of the BGP routing protocol operating among mobile military networks is needed.

To evaluate network routing protocols, the practice of network simulation and emulation is widely used. For example, simulation has been used to perform sensitivity analysis of the OSPF-MDR MANET routing protocol over a broad parameter space [11]. However, simulation results can sometimes be misleading due to modeling and statistical analysis errors [12]. Such modeling discrepancies were seen in an emulation-based evaluation of the OSPF-MDR routing protocol [13]. Here, the simulation's (software) model of the OSPF-MDR routing protocol notably differed from its deployable implementation in terms of its overhead produced. As such, this paper performed its performance evaluation of BGP using a network emulation testbed which executes a deployable version of the BGP routing protocol (e.g., Quagga [23]). Performance data (e.g., routing overhead, end-to-end network reachability, and network outage times) was then derived from the analysis of the networking traffic carried among the real-time components (e.g., Quagga instances) running within the network emulation testbed.

The outline of this paper is as follows. Section II describes previous work evaluating BGP in tactical networks. Section III defines our experimental framework for measuring BGP performance under the stimulus of node mobility in our network emulation testbed. Section IV presents the performance results of our experimentation. Specifically, the measured routing protocol overhead, destination reachability, and network outages times are analyzed and discussed. Section V concludes the paper with discussion of future work.

II. Previous Work

Several works have analyzed the operation and performance of BGP in tactical networks. For example, in [14] the overhead generated by BGP was analyzed for the following: network initialization, steady-state, an increasing number of links failures, and the addition of a new BGP node. However, the authors did not assess the effects of node mobility on BGP.

In [15], network emulation was used to evaluate BGP between nodes within an airborne and SATCOM layer. This previous work varied BGP timer settings while measuring BGP routing convergence time and routing overhead. However, the scale of this emulation experiment was limited to 5 nodes (2 airborne and 3 satellite routers) and the node mobility was restricted to the two airborne nodes. Our work goes further by evaluating BGP in networks of increasing network sizes (up to tens of nodes) and where every node is experiencing some mobility.

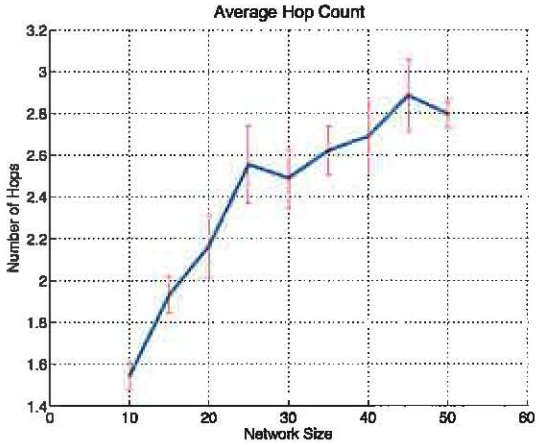

Fig. 1. Average Hop Count across 10 *mobgen-ss* Mobility Realizations

In [16], the authors used simulation to assess BGP's convergence time in response to intermittent link failures, addition of new nodes, and node mobility over various settings of the TCP timers and BGP timers. These authors used a network under test which included 455 BGP routers and where each BGP router was the only member of a unique autonomous system (AS). The authors concluded that the settings of the BGP's hold timer could not improve routing convergence time when both node mobility and intermittent links were present. The authors further advocate the need for dynamic BGP neighbor discovery to lower BGP peering establishment times. Our work provides and evaluates a practical implementation of a dynamic BGP peering mechanism. Additionality, our work evaluates BGP using network emulation and over increasing network sizes.

III. Experimental Framework

A networking experiment's goal is to measure a network's response to some stimulus within some defined networking context. Following experimentation previously presented in [13], this work subjects four routing protocols (including OSPF and BGP) to identical (random) node mobility, limited node transmission ranges, and increasing network sizes. While being exposed to such testing conditions, the routing protocols were evaluated with respect to their generated protocol overhead, their ability to provide valid forwarding paths to

IP addressed destinations (i.e., destination reachability), and the duration of their network's outage events. This experimentation allows for individual performance assessment of each routing protocol as well as comparative performance assessment between multiple routing protocols.

A. Mobile Network Test Topologies

Tactical network topologies are expected to change frequently as node mobility will move nodes in and out of other nodes' transmission ranges. Although commonly criticized, the random waypoint mobility model has become the defacto standard for generating node mobility realizations for simulation and emulation testing [17]. Briefly, under the random node mobility model, each node is assigned a random starting position, followed by instructions to move in a straight line toward some randomly generated waypoint, at some random speed, then pause for some period of time.

There are many implementations of the random waypoint node mobility model. Our work used the *mobgen-ss* mobility model [18] whose nodes' spatial distribution is stationary throughout the generated mobility realization. This minimizes the need for an experimental warm-up period (e.g., initial settling time). Note, the use of a random node mobility generator was chosen due to its simplicity and popularity even though random node motion is not expected to model realistic node mobility associated with tactical military networks.

The *mobgen-ss* mobility generator was used to produce 10 realizations of mobile networks containing 10, 15, ..., 50 nodes. Following [11], each mobile node was specified to have an average speed of 10 m/s, a fixed pause time of 40 secs, operating within a square, bounded area. Furthermore, all *mobgen-ss* mobility realizations were rejected if its network partitioning rate was outside of the range of 5 +/- 0.5% (as reported the SCORES[2] topological analysis tool). This network partitioning constraint was advocated in [19] for the testing of MANET routing protocols. The average (shortest) hop count (as calculated by SCORES) over all 10 *mobgen-ss* realizations for each specific network size is shown in Figure 1.

B. Mobile Nodes as Micro-ASes

Similar to previous work [16], every node captured within a *mobgen-ss* mobility realization will be considered the sole member of a separate, unique autonomous system[3] (AS). Specifically, each *mobgen-ss* mobile node was modeled as a single BGP routing process capable of forming eBGP peering sessions with other mobile nodes. Each such BGP-enabled node can be assumed to be a cluster of connected mobile military units (e.g., strike package, ground patrol). This modeling abstraction is called a *micro-AS*. The use of micro-ASes provides a modeling convenience by removing many of BGP's many known challenges in mobile networks including AS splits and merges [20]), routing domain boundary and gateway definition [21], as well as by eliminating any potential for path loops among internal BGP (iBGP) peers [22].

C. Network Emulation Testbed

To evaluate a routing protocol running among micro-ASes (i.e., mobile networks), a network emulation testbed composed of tens of Linux-based workstations was used. Specifically, the IP networking protocols (e.g., IPv6, TCP, BGP, OSPF) for each micro-AS node was modeled using a single "bare-metal" RHEL5 Linux workstation (i.e., no virtualization is used) and the Quagga software routing suite [23]. Dynamic inter-connectivity between the emulated micro-AS nodes was provided via an Ethernet switch fabric and intelligent testbed control software. Specifically, all outgoing networking traffic from an emulated micro-AS node was broadcast to all other micro-AS nodes in the network emulation experiment. In parallel, the emulation testbed also broadcasts every micro-AS's virtual geographic location over a separate, out-of-band control channel. At each emulated micro-AS node, a local calculation is made to determine which other micro-AS nodes are within its wireless range. Based on these local range calculations, any broadcast emulation traffic arriving at any emulated micro-AS node that is considered out of range (e.g., over 250 meters) is discarded. Such traffic filtering allows for modeling of a tactical network's partial mesh connectivity (i.e., not every node is directly connected to every other node) induced by the micro-AS's finite wireless transmission range as well as the dynamic link interconnectivity induced by micro-AS node mobility.

D. DynamicBGP: BGP for Mobile Micro-ASes

Typically, BGP routers are configured to connect to predefined neighbors using operator-supplied information in local router configuration files. In the tactical environment, node mobility will affect which nodes are routing neighbors of each other such that static configuration files are not applicable. To dynamically manage a node's BGP neighbors (i.e., its BGP peering sessions), several commercial implementations of BGP peering management were evaluated. These included Cisco's Dynamic BGP Neighbors [24], Juniper's BGP Dynamic Peering [25], and TSAT's Dynamic BGP Peering Process [26]. However, all these techniques were found to require static configuration of at least one BGP node involved in the BGP peering session. As such, these techniques were considered of limited utility for mobile networks.

Instead, the IPv6 neighbor discovery[4] (ND) protocol [27] was integrated with the BGP protocol as shown in Figure 2. Here, the IPv6 ND protocol was used to discover which other neighboring BGP routers become present (or absent) using an advertisement/response handshake mechanism. The IPv6 ND protocol was also capable of determining every neighboring router's IPv6 link-layer address. This information was then used to dynamically establish or tear-down each micro-ASes

[2]http://toilers.mines.edu/MANET-Simulations/scores-shtml

[3]An autonomous system (AS) is a separate administered routing domian.

[4]Note, although a MANET Neighbor Discovery protocol does exists [28], IPv6's neighbor discovery protocol was chosen due to its availablility in the Quagga routing suite.

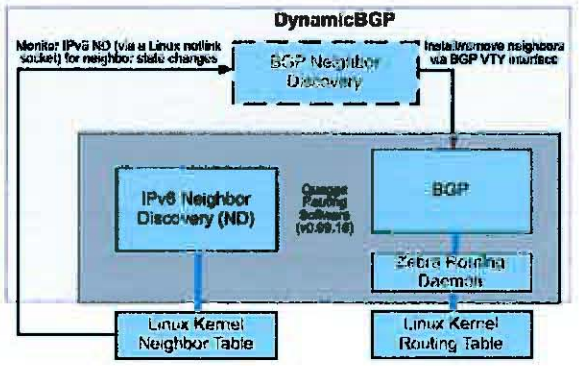

Fig. 2. DynamicBGP Architecture

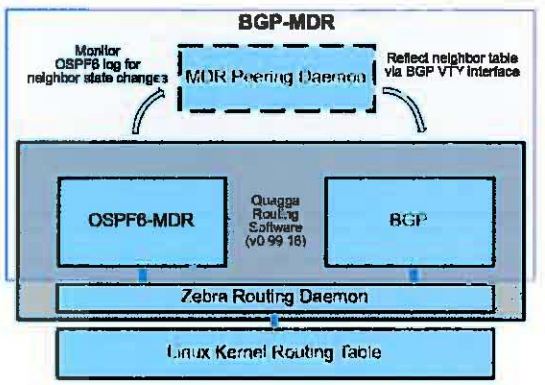

Fig. 3. BGP-MDR Architecture

BGP peering sessions based on the geographical relationships of the micro-ASes (i.e., in- or out-of wireless TX range).

E. BGP-MDR: BGP over a Connected Dominating Set

As discussed above, it is known that BGP is not an overhead efficient routing protocol. This is due in part to BGP's use of path vectors and TCP. Fortunately, several MANET routing protocols (e.g., OSPF-MDR[29]) have shown that the amount of routing information (i.e., overhead) generated in a mobile network can be reduced by using a connected dominating set (CDS). A CDS is a maximum leaf spanning tree induced on the mobile network to provide a backbone over which all routing nodes can communicate. Here, the benefits of a CDS include decreasing number of routing messages (e.g., OSPF LSAs) being flooded by the routing protocol and in reducing the number of routing adjacencies maintained by the mobile nodes. This concept was extended to BGP resulting in the BGP-MDR protocol. The architecture of BGP-MDR within a micro-AS is shown in Figure 3. Here, the (bi-connected) CDS algorithm was executed within an OSPF-MDR protocol instance (running on a separate wireless interface). Then, the routing adjacencies determined by the OSPF-MDR CDS calculation were extracted and used to dynamically configure BGP's peering neighbors. It is noted that this BGP-MDR implementation is only an approximation, as the communication and CPU overhead associated with the CDS algorithm is not occurring with the BGP routing process.

F. Experimental Measurements

To evaluate the performance of several routing protocols (OSPF, BGP) operating within the emulation testbed, two measurement tools were used. The first was *wireshark*[5] which recorded traces of network traffic crossing over the emulation testbed's Ethernet switch fabric. These network traffic traces were then analyzed to derive the amount of network overhead generated by each routing protocol under test.

The second tool uses active ping probes to develop a "pingmatrix". Specifically, for a network containing N micro-ASes, each micro-AS node ran *fping*[6] once every second. The return responses from every *fping* invocation were then compiled into a NxN "pingmatrix" like the ones shown in left-side panels of Figure 4. The rows and columns of each "pingmatrix" reference specific source and destination nodes respectively, whereas the individual "pingmatrix" cells indicidate the end-to-end connectivity state between the cell's associated source and destination node. More generally, for each "pingmatrix" entry, a green cell indicates the round-trip time (in usec) of an active ping probe between the cell's associated source (i.e., row) and destination (i.e., column) node. A red cell indicates the number of measurement intervals (e.g., seconds) a return response has gone missing between the cell's source and destination node. Every red cell represents an active network outage event (i.e., some network node pair cannot communicate).

From these "pingmatrix" semantics, instantaneous destination reachability can be defined as the ratio of the current number of green cells (i.e., returned ping responses) to the total number of "pingmatrix" cells (i.e., N^2). Clearly, the average destination reachability is calculated over all instantaneous destination reachability measurements for the experiment. With respect to network outage events, the total number of outage events is determined by the number of green to red cell transitions. As the duration of each outage event is the number of measurement intervals between when the outage event started (i.e., some green cell turns red) and when the outage event ends (i.e., the red cell turns green), the average network outage time is taken as the average over all network outage event times. These "pingmatrix" derived metrics were both stored in a database for post-experiment analysis as well as made available in real-time during the experiment (see the upper-middle "Plotter" panel of Figure 4).

G. Other Experimental Configuration

Experiments to evaluate the performance of a routing protocol within a mobile network generally contain many con-

[5]http://www.wireshark.org

[6]*fping* is a network measurement application that sends/collects ping probes to/from multiple destinations. See http://fping.sourceforge.net

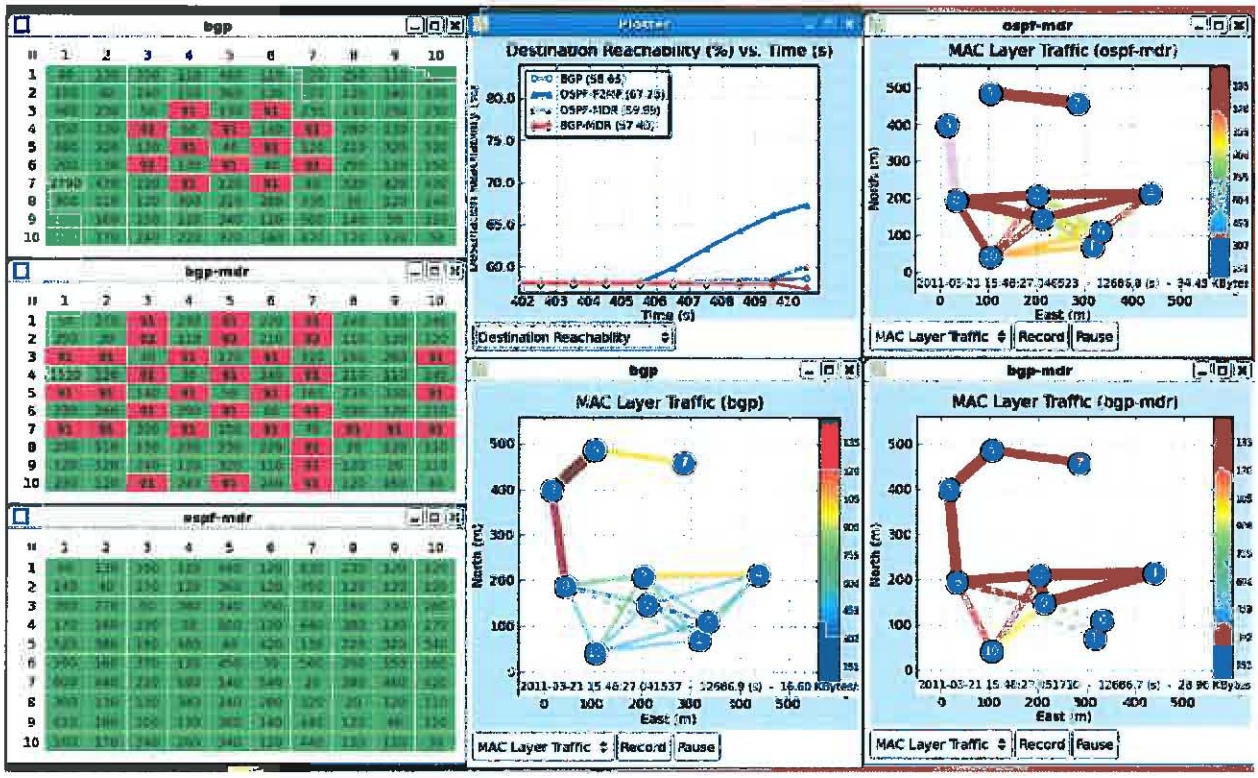

Fig. 4. Live measurement snapshot of a running emulation experiment after a network join event. Specifically, at around time 404s, node 8 moves within node 3's transmission range of 250m allowing the cluster formed by nodes 3, 5, and 7 to rejoin the rest of the network. For several 10 node network instances under test, the left-hand "pingmatrix" panels show the current (i.e., time 411s) instantaneous network reachability between all source/destination pairs. The "pingmatrix" green cells indicate the end-to-end latency (in usec) reported by active pings probes. The red cells indicate the number of measurement intervals (e.g., secs) since the last valid ping response was received (i.e., current outage duration). The center "Plotter" panel shows the average destination reachability for all multiple routing protocols currently under test. The remaining three panels show the nodes' positions and inter-node traffic rate for (3 out of 4) network instances executing some routing protocol under test.

Parameter	Setting
Node Speed	10 +/- 9.99 m/s
Node Pause Time	40 s
Node Location Update Frequency	1 s
Node TX Range	250 m
Node Wireless Bandwidth	10Mbps
Channel Model	100% or 0% loss based on TX range
OSPF Hello Timer	2 s
OSPF Dead Timer	6 s
OSPF Min LSA Adv. Timer	0 s
OSPF-MDR Adjacencies	Bi-connected
BGP Hold Timer	2 s
BGP Keep-Alive Timer	6 s
BGP MRAI Timer	3 s
No. of Mobility Files per Network Size	10
Mobility File Length	1500 s
Node Position Update Rate	1 s
Network Partitioning Rate	5%

TABLE I
EXPERIMENTATION CONFIGURATION PARAMETERS

figuration parameters. Space precludes listing all parameters used in our experiments. However, Table I is provided to capture those parameters deemed most important to understand this work's performance analysis.

IV. RESULTS AND DISCUSSION

As discussed in Section III-F, two tools were used to collect measurement data from the network emulation testbed: *wireshark* and the "pingmatrix". The measurement data from the former was parsed to isolate the overhead produced by each routing protocol under test, whereas the latter was used to determine network reachability (i.e., is there at least one valid network path between every network node pair) and network outage times (i.e., how long were some network node pairs disconected). The performance metrics derived from these two measurement tools is presented and discussed in the following sections.

A. Overhead Analysis

Figure 5 plots the time-averaged protocol overhead generated by several different routing protocols for mobile networks of increasing size. Its legend labels OSPF-MDR and OSPFv3-PtoMpt refer to data calculated for those network instances in which the IETF's OSPFv3 protocol was running with either its wireless interface configured as either the MANET or point-to-multipoint type, respectively. Similarly, Figure 5 legend labels BGP and BGP-MDR refer to network instances running each of this work's modified versions of BGP (both were presented in Section III-D and III-E respectively.) From Figure 5, it is

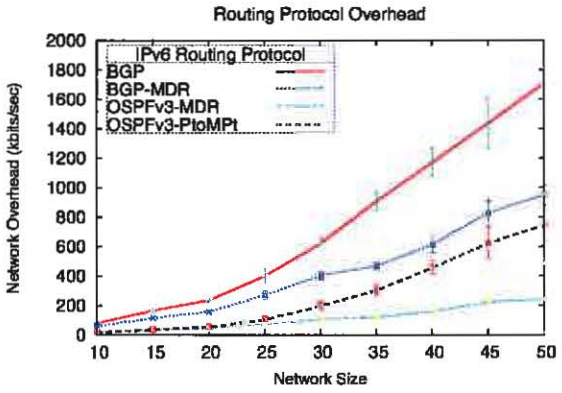

Fig. 5. Time-Averaged Routing Overhead over 10 Mobility Realizations

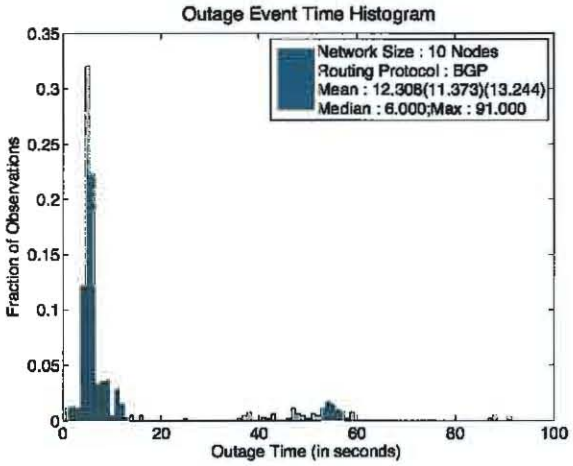

Fig. 6. Sample Outage Time Distribution. 10 Micro-Ases running the BGP Routing Protocol)

seen that the amount of BGP overhead can be several times higher[7] than the dedicated MANET routing protocol (e.g., OSPF-MDR). Specifically, for our experiments in which the network under test contained 50 micro-ASes, BGP averaged about 1600 kbits/s whereas OSPF-MDR averaged about 200 kbits/sec (for an 8x increase). Similarly, BGP also produced notably more overhead than OSPF-PtoMPt (i.e., 1600 kbits/sec versus 750 kbits/sec). Note also that BGP-MDR did notably decrease the amount of BGP overhead (i.e., 1600 kbits/sec down to 850 kbits/sec) indicating that the overhead reduction benefit of a connected dominating set does extend to path vector routing protocols such as BGP. Lastly, by evaluating the BGP overhead with respect to particular protocol message types (e.g., BGP OPEN, BGP UPDATES), it was seen that TCP-related messages and BGP UPDATE messages account for approximately 20% and 50% of the total BGP overhead (not shown due to space constraints).

B. Network Outages

The BGP routing protocol is known to converge slowly, especially after a some destination prefix is withdrawn [5]). Through evaluation of the "pingmatrix" outage events, an estimate of BGP's relative convergence time was developed for our mobile networks. Specifically, the distribution of the "pingmatrix" outage times was generated for every routing protocol under test. In Figure 6, the outage time distribution for BGP is shown. Although Figure 6 shows many outage events whose duration is around 5 secs, the distribution clearly contains several outliers and is not normal. Further analysis using the distribution's average is likely not informative. Instead, the median path outage times were extracted from each routing protocol's outage time distributions. These median outage times are shown in Figure 7. For larger network sizes, the BGP and BGP-MDR routing protocols are seen to result in longer median outage times than OSPF-MDR and OSPF-PtoMPt, but the time difference is only on the order of seconds.

However, an analysis of the maximum outage times for each routing protocol shows that BGP will generally result in outage events that are notably longer than OSPF. Specifically, in Figure 8, BGP (or BGP-MDR) is seen to produce maximum outages times that are tens of seconds[8] longer than OSPF (or OSPF-MDR). Furthermore, in Figure 9 it is seen that a network running OSPF or OSPF-MDR generally has more outage events than if the network runs BGP or BGP-MDR. It is claimed that the OSPF protocol variations converge more quickly, leading to shorter outage events, whereas the BGP protocol variations reacts more slowly, and tends to combined multiple network outage events together into fewer, but longer, outage events. However, this combination effect appears infrequent with respect to very long outage events (i.e., those on the order of tens of seconds). For example, in the BGP outage time distribution of Figure 6), the fraction of outage event times greater than 60 sec is seen to be small. It is believed that the low occurrence rate of very long outage events (i.e., greater than 60 secs) is influenced by the lower number of hops in the networks under test (see Figure 1). To confirm this claim, future work will analyze the lifetime of the entries within each protocol's routing table as well as performing experimentation in networks with longer average shortest path lengths.

C. Reachability Analysis

The average number of node pairs connected by the routing protocols under test is shown in Figure 10. Although each protocol generated different amounts of overhead (see Figure 5), the average destination reachability for all routing protocols was similar. This is likely due to the median outage rate

[7]Note that our comparison of OSPF to BGP may be inaccurate due to different route advertisement timer settings. See the BGP MRAI and OSPF LSA Adv(ertisement) timer settings in Table I. Future work will evaluate this sensitivity.

[8]Note the overall duration of most outage events is highly influenced by the experiment's mobility realization which contains at least 5% network partitioning.

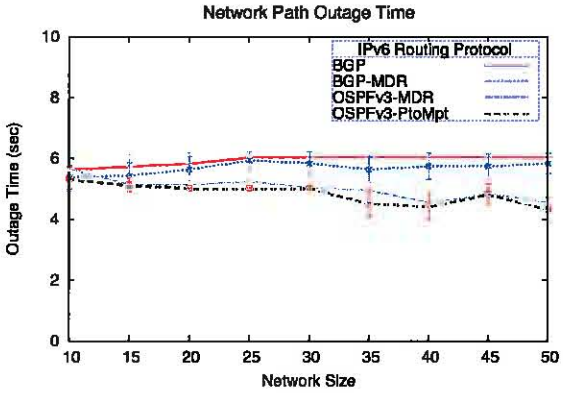

Fig. 7. Average Median Network Outage Times over 10 Mobility Realizations

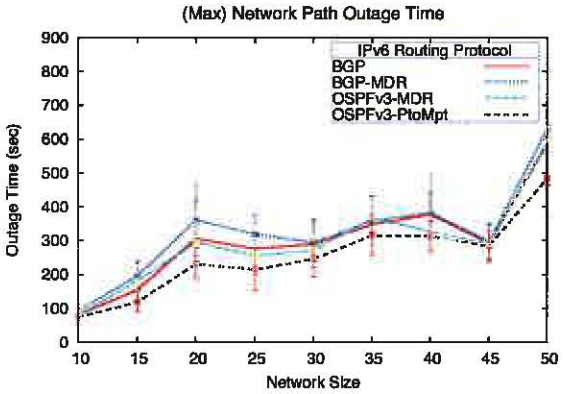

Fig. 8. Average Maximum Outage Times over 10 Mobility Realizations

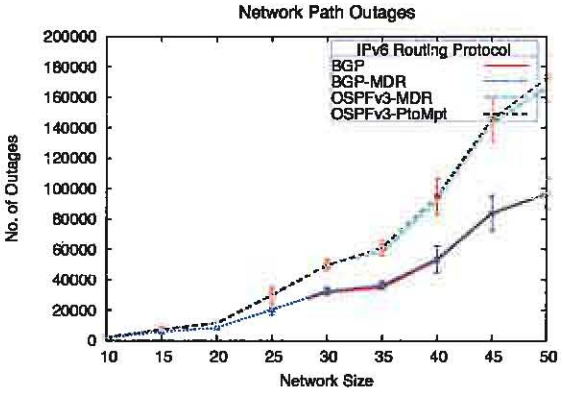

Fig. 9. Average Number of Network Outage Events over 10 Mobility Realizations

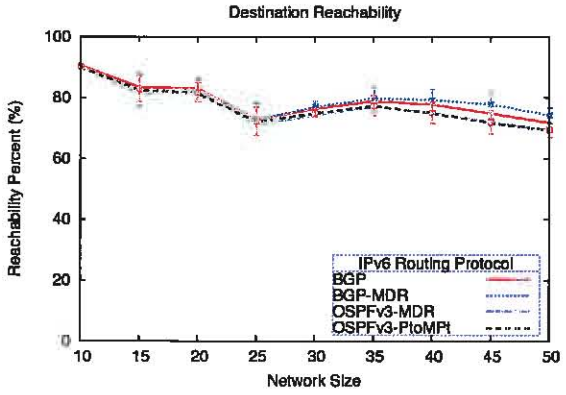

Fig. 10. Average Destination Reachability over 10 Mobility Realizations

of each routing protocol (see Figure 7) being similar and the infrequency of very long outage events (see Figure 6). However, as shown in Figure 4, the instantaneous reachability associated with each protocol can differ (i.e., at time 411s, OSPF is at 67% whereas BGP is at 58%). Future work will perform additional experimentation and analysis to better describe the interactions behind this reachability performance result.

V. Conclusions and Future Work

The motivation for this work was to determine whether the de-facto policy-based routing protocol BGP should be further explored to help manage the increased interconnectivity envisioned between Service and Coalition networks. To this end, this paper described a performance assessment of the BGP routing protocol operating between several emulated mobile networks. Our experimental results appear to confirm several performance limitations of BGP: 1) BGP has high overhead and 2) BGP converges slowly. However, our work also showed that both limitations are either able to be mitigated or occur less frequently than assumed for one class of networks under test (e.g., networks with a 5% partitioning range, nodes with a TX range of 250m, average shortest path lengths less than 3 hops)

For example, our performance testing showed that although BGP scales poorly with respect to its generated overhead, techniques that have proven useful in many MANET routing protocols (e.g., OSPF-MDR) can be similarly used to reduce the amount of protocol overhead. Specifically, this paper evaluated a modified version of BGP (called BGP-MDR) which used a connected dominating set (CDS) algorithm to form a routing backbone overlay network. Notable reduction in BGP's generated overhead was seen with only a minor variation in average destination reachability and median network outage time. Future overhead reduction techniques that can be considered for BGP will include wireless TCP enhancements (e.g., [30]) and adapting BGP to use a node's wireless interface broadcast capability (instead of point-to-point).

With respect to BGP's slow convergence property, our analysis of network outage events showed that for networks running BGP instead of OSPF, fewer outage events will occur, but the median duration of those events will be similar. However, BGP (and BGP-MDR) will produce outage events that are tens of seconds beyond that encountered by OSPF (and OSPF-MDR). However, such long outage times were also seen to be relatively infrequent. This claim will be further explored via analyzing the lifetime of routes (rather than inference via forwarding path pings). If it is confirmed that BGP's longer outage times are due to path hunting [5], all previous techniques to improving BGP's path exploration efficiency (e.g., BGP-RCN [31]) can be explored.

Lastly, we note that BGP has other operational concerns that were not addressed in this work. For example, BGP has the ability to cause wide-scale network disruptions (e.g., YouTube hijacking by Pakistan Telecom [6]) due to the protocol's configuration complexity [7] and its security weaknesses [8]. Furthermore, BGP can become unstable (i.e., oscillate) based on certain configurations of its routing polices [9], [10]. Such issues will be investigated in future work.

In summary, several of BGP's challenges within mobile networks are somewhat similar to those previously seen for OSPF. As such, there exists opportunity to build on ideas and lessons learned by adapting OSPF to the tactical edge. It is thought that such ideas can combined with other known BGP efficiency improvements to produce a policy-based routing protocol for tactical networks. If an efficient policy-based BGP routing protocol can be made available for the tactical edge, then the evaluation of policy-based management between Service and Coalition networks can begin.

References

[1] JTRS Joint Program Office, "Joint tactical radio system (JTRS) wide-band networking waveform (WNW) functional description document (FDD)," November 21 2001.

[2] Y. Reknter, T. Li, and S. Haren, "A border gateway protocol 4 (BGP-4)," http://www.rfc-archive.org/getrfc.php?rfc=4271, January 2006.

[3] C.-K. Toh, Ad Hoc Mobile Wireless Network: Protocols and Systems. Prentice Hall, 2002.

[4] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "Observation and analysis of BGP behavior under stress," in IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment. New York, NY, USA: ACM Press, 2002, pp. 183–195.

[5] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," in SIGCOMM, Stockholm, Sweden, August 2000, pp. 175–187. [Online]. Available: http://citeseer.ist.psu.edu/labovitz00delayed.html

[6] R. Blog, "Pakistan hijacks YouTube," http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml.

[7] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in ACM SIGCOMM, Pittsburgh, August 2002.

[8] K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A survey of BGP security," Pennsylvania State University, http://www.patrickmcdaniel.org/pubs/td-5ugj33.pd, Tech. Rep., 2005.

[9] K. Varadhan, R. Govindan, and D. Estrin, "Persistent route oscillations in interdomain routing," Computer Networks, vol. 32, no. 1, pp. 1–16, 2000. [Online]. Available: http://citeseer.ist.psu.edu/article/varadhan99persistent.html

[10] T. Griffin, F. B. Shepard, and G. Wilfong, "The stable paths problem and interdomain routing," IEEE/ACM Transactions on Networking, vol. 10, no. 2, pp. 232–243, April 2002.
[11] T. R. Henderson, P. A. Spagnolo, and G. Pei, "Evaluation of OSPF MANET extensions," The Boeing Company, http://hipserver.mct.phantomworks.org/ietf/ospf/reports/Boeing-D950-10897-1.pdf, Tech. Rep. D950-10897-11, Last accessed in March 2011 2005.
[12] S. Kurkowski, T. Camp, and M. Colagrosso, "MANET simulation studies: The incredibles," SIGMOBILE Mob. Comput. Commun. Rev., vol. 9, no. 4, pp. 50–61, 2005.
[13] G. Carl, S. Dastangoo, and J. Wildman, "Using an emulation testbed to measure OSPF routing overhead due to mobility in wireless ad hoc networks," in Military Communications Conference, 2008. MILCOM 2008. IEEE, San Diego, CA, Nov. 2008, pp. 1–9.
[14] M. Gobrial, "Evaluation of border gateway protocol (BGP) version 4 (V4) in the tactical environment," in MILCOM '96, vol. 2, October 1996, pp. 490–495.
[15] K. Schroth and D. Kiwior, "Interdomain routing for mobile nodes," in MILCOM, Orlando, FL, October 2007, pp. pp.1–8.
[16] S. Erramilli, J. Lee, L. Kant, A. McAuley, J. Giacopelli, K. Adams, and J. Pulliam, "Performance modeling and analysis of routing in heterogeneous multi-tier ad-hoc networks," in MILCOM, vol. 3, Atlantic City, NJ, October 2005, pp. 1555–1561.
[17] J. Yoon, M. Liu, and B. Noble, "Random waypoint considered harmful," in INFOCOM '03, 2003.
[18] W. Navidi and T. Camp, "Stationary distributions for the random waypoint mobility model," IEEE Transactions on Mobile Computing, vol. 3, pp. 99–108, 2003.
[19] S. Kurkowski, T. Camp, and W. Navidi, "Two standards for rigorous MANET routing protocol evaluation," in Mobile Adhoc and Sensor Systems (MASS), 2006 IEEE International Conference on, Oct. 2006, pp. 256–266.
[20] C.-K. Chau, J. Crowcroft, K.-W. Lee, and S. H. Wong, "IDRM: Inter-domain routing protocol for mobile ad hoc networks," University of Cambridge, Tech. Rep. UCAM-CL-TR-708, 2008.
[21] P. A. Spagnolo and T. R. Henderson, "Connecting OSPF MANET to larger networks," in MILCOM, Boston, October 2007, pp. 1–7.
[22] T. G. Griffin and G. Wilfong, "On the correctness of IBGP configuration," in SIGCOMM '02: Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications. New York, NY, USA: ACM Press, 2002, pp. 17–29.
[23] K. Ishiquro, "Quagga software routing suite," Available at http://www.quagga.net/ [Last accessed March 29, 2010], March 2010.
[24] "Dynamic BGP neighbors," Cisco IOS Software Releases 12.2 SX.
[25] "BGP dynamic peering," Juniper JUNOSe 8.1.x BGP and MPLS Configuration Guide.
[26] N. Lovering and W. Wheeler, "Dynamic BGP peering process," TSAT Network Architecture Working Group, Tech. Rep. v1.2, November 2004.
[27] T. Narten, E. Nordmark, and W. Simpson, "Neighbor discovery for ip version 6 (ipv6)," http://tools.ietf.org/html/rfc2461 [Last accessed March 29, 2010], December 1998.
[28] T. Clausen, C. Dearlove, and J. Dean, "Mobile ad hoc network (MANET) neighborhood discovery protocol (NHDP)," http://tools.ietf.org/html/draft-ietf-manet-nhdp-12 [Last accessed March 29, 2010], March 2010.
[29] R. Ogier and P. Spagnolo, "MANET extension of OSPF using CDS flooding," http://tools.ietf.org/html/draft-ietf-ospf-manet-mdr-02, June 2008.
[30] A. K. Singh and S. Iyer, "Atcp: Improving tcp performance over mobile wireless environments," in In IEEE ICPWC, 2002, pp. 81–85.
[31] D. Pei, M. Azuma, D. Massey, and L. Zhang, "BGP-RCN: Improving BGP convergence through root cause notification," Comput. Netw. ISDN Syst., vol. 48, no. 2, pp. 175–194, 2005.